

A Gentle Tutorial of Recurrent Neural Network with Error Backpropagation

Gang Chen *

Department of Computer Science and Engineering, SUNY at Buffalo

1 abstract

We describe recurrent neural networks (RNNs), which have attracted great attention on sequential tasks, such as handwriting recognition, speech recognition and image to text. However, compared to general feedforward neural networks, RNNs have feedback loops, which makes it a little hard to understand the backpropagation step. Thus, we focus on basics, especially the error backpropagation to compute gradients with respect to model parameters. Further, we go into detail on how error backpropagation algorithm is applied on long short-term memory (LSTM) by unfolding the memory unit.

2 Sequential data

Sequential data is common in a wide variety of domains including natural language processing, speech recognition and computational biology. In general, it is divided into time series and ordered data structures. As for the time-series data, it changes over time and keeps consistent in the adjacent clips. For example the time frames for speech or video analysis, daily prices of stocks or the rainfall measurements on successive days. There are also ordered data in the sequence, such as text and sentence for handwriting recognition, and genes. For example, successfully predicting protein-protein interactions requires knowledge of the secondary structures of the proteins and semantic analysis might involve annotating tokens with parts of speech tags.

The goal of this work is for sequence labeling, i.e. classify all items in a sequence [4]. For example, in handwritten word recognition we wish to label a sequence of characters given features of the characters; in part of speech tagging, we wish to label a sequence of words.

More specifically, given an observation sequence $\{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_T\}$ and its corresponding label $\{y_1, y_2, \dots, y_T\}$, we want to learn a map $f : \mathbf{x} \mapsto y$. In the following, we will introduce RNNs to model the sequential data and give details on backpropagation over the feedback loop.

*gangchen@buffalo.edu, SUNY at Buffalo




Figure 1: It is a RNN example: the left recursive description for RNNs, and the right is the corresponding extended RNN model in a time sequential manner.

3 Recurrent neural networks

A RNN is a kind of neural networks, which can send feedback signals (to form a directed cycle), such as Hopfield net [1] and long-short term memory (LSTM) [2]. RNN models a dynamic system, where the hidden state \mathbf{h}_t is not only dependent on the current observation \mathbf{x}_t , but also relies on the previous hidden state \mathbf{h}_{t-1} . More specifically, we can represent \mathbf{h}_t as

$$\mathbf{h}_t = f(\mathbf{h}_{t-1}, \mathbf{x}_t) \quad (1)$$

where f is a nonlinear mapping. Thus, \mathbf{h}_t contains information about the whole sequence, which can be inferred from the recursive definition in Eq. 1. In other words, RNN can use the hidden variables as a memory to capture long term information from a sequence.

Suppose that we have the following RNN model, such that

$$\mathbf{h}_t = \tanh(W_{hh}\mathbf{h}_{t-1} + W_{xh}\mathbf{x}_t + \mathbf{b}_h) \quad (2)$$

$$z_t = \text{softmax}(W_{hz}\mathbf{h}_t + \mathbf{b}_z) \quad (3)$$

where z_t is the prediction at the time step t , and $\tanh(x)$ is defined as

$$\tanh(x) = \frac{\sinh(x)}{\cosh(x)} = \frac{e^x - e^{-x}}{e^x + e^{-x}} = \frac{e^{2x} - 1}{e^{2x} + 1}$$

Take the derivative of $\tanh(x)$ w.r.t. x

$$\begin{aligned} \frac{\partial \tanh(x)}{\partial(x)} &= \frac{\partial \frac{\sinh(x)}{\cosh(x)}}{\partial x} \\ &= \frac{\frac{\partial \sinh(x)}{\partial x} \cosh(x) - \sinh(x) \frac{\partial \cosh(x)}{\partial x}}{(\cosh(x))^2} \\ &= \frac{[\cosh(x)]^2 - [\sinh(x)]^2}{(\cosh(x))^2} \\ &= 1 - [\tanh(x)]^2 \end{aligned} \quad (4)$$

More specifically, the RNN model above has one hidden layer as depicted in Fig. 1. Notice that it is very easy to extend the one hidden case into multiple layers, which has been discussed in deep neural network before. Considering the varying length for each sequential data, we also assume the parameters in each time step are the same across the whole sequential analysis. Otherwise it will be hard to compute the gradients. In addition, sharing the weights for any sequential length can generalize the model well. As for sequential labeling, we can use the maximum likelihood to estimate model parameters. In other words, we can minimize the negative log likelihood the objective function

$$\mathcal{L}(\mathbf{x}, \mathbf{y}) = - \sum_t y_t \log z_t \quad (5)$$

In the following, we will use notation \mathcal{L} as the objective function for simplicity. And further we will use $\mathcal{L}(t+1)$ to indicate the output at the time step $t+1$, s.t. $\mathcal{L}(t+1) = -y_{t+1} \log z_{t+1}$. Let's first take the derivative with respect to z_t

$$\frac{\partial \mathcal{L}}{\partial z_t} = -(y_t - z_t) \quad (6)$$

Note the weight W_{hz} is shared across all time sequence, thus we can differentiate to it at each time step and sum all together

$$\frac{\partial \mathcal{L}}{\partial W_{hz}} = \sum_t \frac{\partial \mathcal{L}}{\partial z_t} \frac{\partial z_t}{\partial W_{hz}} \quad (7)$$

Similarly, we can get the gradient w.r.t. bias b_z

$$\frac{\partial \mathcal{L}}{\partial b_z} = \sum_t \frac{\partial \mathcal{L}}{\partial z_t} \frac{\partial z_t}{\partial b_z} \quad (8)$$

Now let's go through the details to derive the gradient w.r.t. W_{hh} . Considering at the time step $t \rightarrow t+1$ in Fig. 1,

$$\frac{\partial \mathcal{L}(t+1)}{\partial W_{hh}} = \frac{\partial \mathcal{L}(t+1)}{\partial z_{t+1}} \frac{\partial z_{t+1}}{\partial \mathbf{h}_{t+1}} \frac{\partial \mathbf{h}_{t+1}}{\partial W_{hh}} \quad (9)$$

where we only consider one step $t \rightarrow (t+1)$. And because the hidden state \mathbf{h}_{t+1} partially depends on \mathbf{h}_t , so we can use backpropagation to compute the above partial derivative. Think further W_{hh} is shared cross the whole time sequence, according to the recursive definition in Eq. 2. Thus, at the time step $(t-1) \rightarrow t$, we can further get the partial derivative w.r.t. W_{hh} as follows

$$\frac{\partial \mathcal{L}(t+1)}{\partial W_{hh}} = \frac{\partial \mathcal{L}(t+1)}{\partial z_{t+1}} \frac{\partial z_{t+1}}{\partial \mathbf{h}_{t+1}} \frac{\partial \mathbf{h}_{t+1}}{\partial \mathbf{h}_t} \frac{\partial \mathbf{h}_t}{\partial W_{hh}} \quad (10)$$

Thus, at the time step $t+1$, we can compute gradient w.r.t. z_{t+1} and further use backpropagation through time (BPTT) from t to 0 to calculate gradient w.r.t. W_{hh} , shown as the red chain in Fig. 1. Thus, if we only consider the output z_{t+1} at the time step $t+1$, we can yield the following gradient w.r.t. W_{hh}

$$\frac{\partial \mathcal{L}(t+1)}{\partial W_{hh}} = \sum_{k=1}^t \frac{\partial \mathcal{L}(t+1)}{\partial z_{t+1}} \frac{\partial z_{t+1}}{\partial \mathbf{h}_{t+1}} \frac{\partial \mathbf{h}_{t+1}}{\partial \mathbf{h}_k} \frac{\partial \mathbf{h}_k}{\partial W_{hh}} \quad (11)$$

Aggregate the gradients w.r.t. W_{hh} over the whole time sequence with back propagation, we can finally yield the following gradient w.r.t. W_{hh}

$$\frac{\partial \mathcal{L}}{\partial W_{hh}} = \sum_t \sum_{k=1}^{t+1} \frac{\partial \mathcal{L}(t+1)}{\partial z_{t+1}} \frac{\partial z_{t+1}}{\partial \mathbf{h}_{t+1}} \frac{\partial \mathbf{h}_{t+1}}{\partial \mathbf{h}_k} \frac{\partial \mathbf{h}_k}{\partial W_{hh}} \quad (12)$$

Now we turn to derive the gradient w.r.t. W_{xh} . Similarly, we consider the time step $t+1$ (only contribution from \mathbf{x}_{t+1}) and calculate the gradient w.r.t. to W_{xh} as follows

$$\frac{\partial \mathcal{L}(t+1)}{\partial W_{xh}} = \frac{\partial \mathcal{L}(t+1)}{\partial \mathbf{h}_{t+1}} \frac{\partial \mathbf{h}_{t+1}}{\partial W_{xh}}$$

Because \mathbf{h}_t and \mathbf{x}_{t+1} both make contribution to \mathbf{h}_{t+1} , we need to backpropagate to \mathbf{h}_t as well. If we consider the contribution from the time step t , we can further get

$$\begin{aligned} \frac{\partial \mathcal{L}(t+1)}{\partial W_{xh}} &= \frac{\partial \mathcal{L}(t+1)}{\partial \mathbf{h}_{t+1}} \frac{\partial \mathbf{h}_{t+1}}{\partial W_{xh}} + \frac{\partial \mathcal{L}(t+1)}{\partial \mathbf{h}_t} \frac{\partial \mathbf{h}_t}{\partial W_{xh}} \\ &= \frac{\partial \mathcal{L}(t+1)}{\partial \mathbf{h}_{t+1}} \frac{\partial \mathbf{h}_{t+1}}{\partial W_{xh}} + \frac{\partial \mathcal{L}(t+1)}{\partial \mathbf{h}_{t+1}} \frac{\partial \mathbf{h}_{t+1}}{\partial \mathbf{h}_t} \frac{\partial \mathbf{h}_t}{\partial W_{xh}} \end{aligned} \quad (13)$$

Thus, summing up all contributions from t to 0 via backpropagation, we can yield the gradient at the time step $t+1$

$$\frac{\partial \mathcal{L}(t+1)}{\partial W_{xh}} = \sum_{k=1}^{t+1} \frac{\partial \mathcal{L}(t+1)}{\partial \mathbf{h}_{t+1}} \frac{\partial \mathbf{h}_{t+1}}{\partial \mathbf{h}_k} \frac{\partial \mathbf{h}_k}{\partial W_{xh}} \quad (14)$$

Further, we can take derivative w.r.t. W_{xh} over the whole sequence as

$$\frac{\partial \mathcal{L}}{\partial W_{xh}} = \sum_t \sum_{k=1}^{t+1} \frac{\partial \mathcal{L}(t+1)}{\partial z_{t+1}} \frac{\partial z_{t+1}}{\partial \mathbf{h}_{t+1}} \frac{\partial \mathbf{h}_{t+1}}{\partial \mathbf{h}_k} \frac{\partial \mathbf{h}_k}{\partial W_{xh}} \quad (15)$$

However, there are gradient vanishing or exploding problems to RNNs. Notice that $\frac{\partial \mathbf{h}_{t+1}}{\partial \mathbf{h}_k}$ in Eq. 15 indicates matrix multiplication over the sequence. Because RNNs need to backpropagate gradients over a long sequence (with small values in the matrix multiplication), gradient value will shrink layer over layer, and eventually vanish after a few time steps. Thus, the states that are far away from the current time step does not contribute to the parameters' gradient computing (or parameters that RNNs is learning). Another direction is the gradient exploding, which attributed to large values in matrix multiplication.

Considering the weakness of RNNs, long short term memory (LSTM) was proposed to handle gradient vanishing problem [2]. Notice that RNNs makes use of a simple *tanh* function to incorporate the correlation between \mathbf{x}_t and \mathbf{h}_{t-1} and \mathbf{h}_t , while LSTM model such correlation with a memory unit. And LSTM has attracted great attention for time series data recently and yielded significant improvement over RNNs. For example, LSTM has demonstrated very promising results on handwriting recognition task. In the following part, we will introduce LSTM, which introduces memory cells for the hidden states.

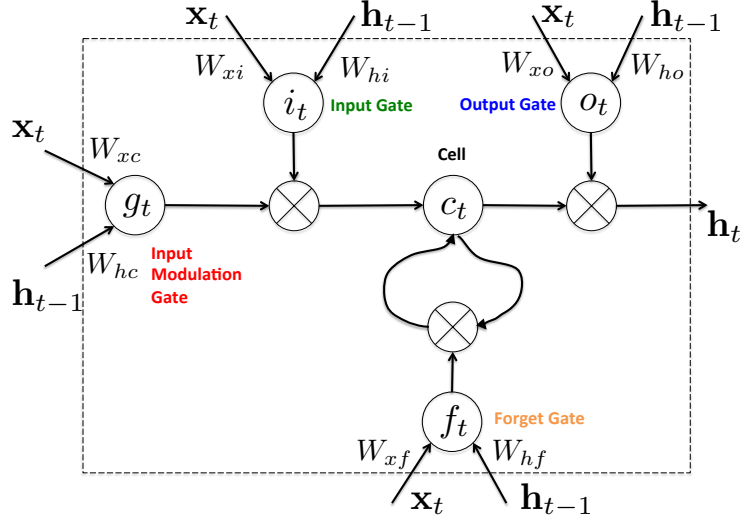


Figure 2: It is an unit structure of LSTM, including 4 gates: input modulation gate, input gate, forget gate and output gate.

3.1 Long short-term memory (LSTM)

The architecture of RNNs have cycles incorporating the activations from previous trim steps as input to the network to make a decision for the current input, which makes RNNs better suited for sequential labeling tasks. However, one vital problem of RNNs is the gradient vanishing problem. LSTM [2] extends the RNNs model with two advantages: (1) introduce the memory information (or cell) (2) handle long sequential better, considering the gradient vanishing problem, refer to Fig. 2 for the unit structure. In this part, we will introduce how to forward and backward LSTM neural network, and we will derive gradients and backpropagate error in details.

The core of LSTM is a memory unit (or cell) \mathbf{c}_t in Fig. 2, which encodes the information of the inputs that have been observed up to that step. The memory cell \mathbf{c}_t has the same inputs (\mathbf{h}_{t-1} and \mathbf{x}_t) and outputs \mathbf{h}_t as a normal recurrent network, but has more gating units which control the information flow. The input gate and output gate respectively control the information input to the memory unit and the information output from the unit. More specifically, the output \mathbf{h}_t of the LSTM cell can be shut off via the output gate.

As to the memory cell itself, it is also controlled with a forget gate, which can reset the memory unit with a sigmoid function. More specifically, given a sequence data $\{\mathbf{x}_1, \dots, \mathbf{x}_T\}$ we have the gate




Figure 3: This graph unfolds the memory unit of LSTM, with the purpose to make it easy to understand error backpropagation.

definition as follows:

$$\mathbf{f}_t = \sigma(W_{xf}\mathbf{x}_t + W_{hf}\mathbf{h}_{t-1} + b_f) \quad (16)$$

$$\mathbf{i}_t = \sigma(W_{xi}\mathbf{x}_t + W_{hi}\mathbf{h}_{t-1} + b_i) \quad (17)$$

$$\mathbf{o}_t = \tanh(W_{xo}\mathbf{x}_t + W_{ho}\mathbf{h}_{t-1} + b_o) \quad (18)$$

$$\mathbf{c}_t = \mathbf{f}_t \circ \mathbf{c}_{t-1} + \mathbf{i}_t \circ \mathbf{g}_t \quad (19)$$

$$\mathbf{g}_t = \sigma(W_{xc}\mathbf{x}_t + W_{hc}\mathbf{h}_{t-1} + b_c) \quad (20)$$

$$\mathbf{h}_t = \mathbf{o}_t \circ \tanh(\mathbf{c}_t), \quad z_t = \text{softmax}(W_{hz}\mathbf{h}_t + b_z) \quad (21)$$

where \mathbf{f}_t indicates forget gate, \mathbf{i}_t input gate, \mathbf{o}_t output gate and \mathbf{g}_t input modulation gate. Note that the memory unit models much more information than RNNs, except Eq. 18.

Same as RNNs to make prediction, we can add a linear model over the hidden state \mathbf{h}_t , and output the likelihood with softmax function

$$z_t = \text{softmax}(W_{hz}\mathbf{h}_t + b_z)$$

If the groundtruth at time t is y_t , we can consider minimizing least square $\frac{1}{2}(y_t - z_t)^2$ or cross entropy to estimate model parameters. Thus, for the top layer classification with weight W_{hz} , we can take derivative w.r.t. z_t and W_{hz} respectively

$$dz_t = y_t - z_t \quad (22)$$

$$dW_{hz} = \sum_t \mathbf{h}_t dz_t \quad (23)$$

$$d\mathbf{h}_T = W_{hz} dz_T \quad (24)$$

where we only consider the gradient w.r.t. \mathbf{h}_T at the last time step T . For any time step t , its gradient will be a little different, which will be introduced later (refer to Eq. 34).

then backpropagate the LSTM at the current time step t

$$d\mathbf{o}_t = \tanh(\mathbf{c}_t)d\mathbf{h}_t \quad (25)$$

$$d\mathbf{c}_t = (1 - \tanh(\mathbf{c}_t)^2)\mathbf{o}_td\mathbf{h}_t \quad (26)$$

$$d\mathbf{f}_t = \mathbf{c}_{t-1}d\mathbf{c}_t \quad (27)$$

$$d\mathbf{c}_{t-1} = \mathbf{f}_t \circ d\mathbf{c}_t \quad (28)$$

$$d\mathbf{i}_t = \mathbf{g}_td\mathbf{c}_t \quad (29)$$

$$d\mathbf{g}_t = \mathbf{i}_td\mathbf{c}_t \quad (30)$$

where the gradient w.r.t. $\tanh(\mathbf{c}_t)$ in Eq. 26 can be derived according to Eq. 4.

further, back-propagate activation functions over the whole sequence

$$\begin{aligned} dW_{xo} &= \sum_t (1 - \mathbf{o}_t^2)\mathbf{x}_td\mathbf{o}_t \\ dW_{xi} &= \sum_t \mathbf{i}_t(1 - \mathbf{i}_t)\mathbf{x}_td\mathbf{i}_t \\ dW_{xf} &= \sum_t \mathbf{f}_t(1 - \mathbf{f}_t)\mathbf{x}_td\mathbf{f}_t \\ dW_{xc} &= \sum_t \mathbf{g}_t(1 - \mathbf{g}_t)\mathbf{x}_td\mathbf{g}_t \end{aligned} \quad (31)$$

Note that the weights W_{xo} , W_{xi} , W_{xf} and W_{xc} are shared across the whole sequence, thus we need to take the same summation over t as RNNs in Eq. 7. Similarly, we have

$$\begin{aligned} dW_{ho} &= \sum_t (1 - \mathbf{o}_t^2)\mathbf{h}_{t-1}d\mathbf{o}_t \\ dW_{hi} &= \sum_t \mathbf{i}_t(1 - \mathbf{i}_t)\mathbf{h}_{t-1}d\mathbf{i}_t \\ dW_{hf} &= \sum_t \mathbf{f}_t(1 - \mathbf{f}_t)\mathbf{h}_{t-1}d\mathbf{f}_t \\ dW_{hc} &= \sum_t \mathbf{g}_t(1 - \mathbf{g}_t)\mathbf{h}_{t-1}d\mathbf{g}_t \end{aligned} \quad (32)$$

and corresponding hiddens at the current time step $t - 1$

$$\begin{aligned} d\mathbf{h}_{t-1} &= (1 - \mathbf{o}_t^2)W_{ho}d\mathbf{o}_t + \mathbf{i}_t(1 - \mathbf{i}_t)W_{hi}d\mathbf{i}_t \\ &\quad + \mathbf{f}_t(1 - \mathbf{f}_t)W_{hf}d\mathbf{f}_t + \mathbf{g}_t(1 - \mathbf{g}_t)W_{hc}d\mathbf{g}_t \end{aligned} \quad (33)$$

$$d\mathbf{h}_{t-1} = d\mathbf{h}_{k-1} + W_{hz}dz_{t-1} \quad (34)$$

where we consider two sources to derive $d\mathbf{h}_{t-1}$, one is from activation function in Eq. 33 and the other is from the objective function at the time step t in Eq. 24.

3.2 Error backpropagation

In this part, we will give detail information on how to derive gradients, especially Eq. 28. As we talked about RNNs before, we can take the same strategy to unfold the memory unit, shown in Fig. 3. Suppose we have the least square objective function

$$\mathcal{L}(\mathbf{x}, \theta) = \min \sum_t \frac{1}{2} (y_t - z_t)^2 \quad (35)$$

where $\theta = \{W_{hz}, W_{xo}, W_{xi}, W_{xf}, W_{xc}, W_{ho}, W_{hi}, W_{hf}, W_{hc}\}$ with biases ignored. To make it easy to understand in the following, we use $\mathcal{L}(t) = \frac{1}{2}(y_t - z_t)^2$.

At the time step T , we take derivative w.r.t. \mathbf{c}_T

$$\frac{\partial \mathcal{L}(T)}{\partial \mathbf{c}_T} = \frac{\partial \mathcal{L}(T)}{\partial \mathbf{h}_T} \frac{\partial \mathbf{h}_T}{\partial \mathbf{c}_T} \quad (36)$$

At the time step $T - 1$, we take derivative of $\mathcal{L}(T - 1)$ w.r.t. \mathbf{c}_{T-1} as

$$\frac{\partial \mathcal{L}(T - 1)}{\partial \mathbf{c}_{T-1}} = \frac{\partial \mathcal{L}(T - 1)}{\partial \mathbf{h}_{T-1}} \frac{\partial \mathbf{h}_{T-1}}{\partial \mathbf{c}_{T-1}} \quad (37)$$

However, according to Fig. 3, the error is not only backpropagated via $\mathcal{L}(T - 1)$, but also from \mathbf{c}_T , thus the final gradient w.r.t. \mathbf{c}_{T-1}

$$\begin{aligned} \frac{\partial \mathcal{L}(T - 1)}{\partial \mathbf{c}_{T-1}} &= \frac{\partial \mathcal{L}(T - 1)}{\partial \mathbf{c}_{T-1}} + \frac{\partial \mathcal{L}(T)}{\partial \mathbf{c}_{T-1}} \\ \frac{\partial \mathcal{L}(T - 1)}{\partial \mathbf{c}_{T-1}} &= \frac{\partial \mathcal{L}(T - 1)}{\partial \mathbf{h}_{T-1}} \frac{\partial \mathbf{h}_{T-1}}{\partial \mathbf{c}_{T-1}} + \frac{\partial \mathcal{L}(T)}{\partial \mathbf{h}_T} \frac{\partial \mathbf{h}_T}{\partial \mathbf{c}_T} \frac{\partial \mathbf{c}_T}{\partial \mathbf{c}_{T-1}} \end{aligned} \quad (38)$$

where we use the chain rule in Eq. 38. Further, we can rewrite Eq. 38 as

$$d\mathbf{c}_{T-1} = d\mathbf{c}_{T-1} + \mathbf{f}_T \circ d\mathbf{c}_T \quad (39)$$

In a similar manner, we can derive Eq. 28 at any time step.

3.3 Parameters learning

Forward: we can use Eqs. 16-21 to update states as the feedforward neural network from the time step 1 to T .

Backward: we can backpropagate the error from T to 1 via Eqs. 25-34. After we get gradients using backpropagation, the model θ can be learnt with gradient based methods, such as stochastic gradient descent and L-BFGS). If we use stochastic gradient descent (SGD) to update $\theta = \{W_{hz}, W_{xo}, W_{xi}, W_{xf}, W_{xc}, W_{ho}, W_{hi}, W_{hf}, W_{hc}\}$, then we have

$$\theta = \theta + \eta d\theta \quad (40)$$

where η is the learning rate. Notice that more tricks can be applied here.

4 Applications

RNNs can be used to handle sequential data, such as speech recognition. RNNs can also be extended into multiple layers in a bi-directional manner. Moreover, RNNs can be combined with other neural networks [3], such as convolutional neural network to handle video to text problem, as well as combining two LSTMs for machine translation.

References

- [1] Hopfield, J. J., Neural Networks and Physical Systems with Emergent Collective Computational Abilities, In Neurocomputing: Foundations of Research, 1988.
- [2] Hochreiter, Sepp and Schmidhuber, Jürgen, Long Short-Term Memory, In Neural Comput., 1997.
- [3] G. Chen and S. N. Srihari, Generalized K-fan Multimodal Deep Model with Shared Representations, 2015.
- [4] Gang Chen, Machine Learning: Basics, Models and Trends, 2016